arXiv:1511.00634v3 [stat.ME] 12 Jul 2016

A Simple and Adaptive Dispersion Regression Model for Count Data

Hadeel S. KALKTAWI, Keming YU and Veronica VINCIOTTI
Department of Mathematics,
College of Engineering, Design and Physical Sciences
Brunel University London, Uxbridge UB8 3PH, U.K.
keming.yu@brunel.ac.uk.

Abstract

Regression for count data is widely performed by models such as Poisson, negative Binomial and zero-inflated regression models. A challenge often faced by practitioners is the selection of the right model to take into account dispersion and excessive zeros, both of which typically occur in count data sets. It is highly desirable to have a unified model that can automatically adapt to the underlying dispersion and that can be easily implemented in practice. In this paper, a discrete Weibull regression model is shown to be able to adapt to different types of dispersions in a simple way. In particular, this model has the ability to model data that are both over and under dispersed relative to Poisson. Additionally, this simple model is shown to be capable of modelling highly skewed count data with excessive zeros, without the need for introducing zero-inflated and hurdle components. Maximum likelihood can be used for efficient parameter estimation. Both simulated examples and real data analyses illustrate the model and its adaptation to modelling over-dispersion, under-dispersion and excessive zeros, in comparison to classical regression approaches for count data.

keywords: Discrete Weibull; Count data; Dispersion; Generalised linear models

1 Introduction

Count data, which refers to the number of times an item or an event occurs within a fixed period of time, commonly arises in many fields. Indeed, examples of count

data include the number of heart attacks or the number of hospitalisation days in medical studies, the number of students absent during a period of time in education studies or the number of times parents perpetrate domestic violence against their child in social science investigations. There is now a great deal of interest in the literature on investigating the relationship between a count response variable and other variables; for example, how the education level of parents can affect the incidence of domestic violence against their children. Methods to address these questions fall in the general area of regression analysis of count data (see Cameron and Trivedi (2013) and Hilbe (2014) among others).

Classical regression models for count data belong to the family of generalised linear models (Nelder and Wedderburn, 1972), such as Poisson regression, which models the conditional mean of the counts as a linear regression on a set of covariates through the log link function. Although Poisson regression is fundamental to the regression analysis of count data, it is often of limited use for real data, due to its property of equal mean and variance. Real data usually experience over-dispersion relative to Poisson or the opposite case of under-dispersion. Thus, accounting for over-dispersion and under-dispersion when modelling count data is essential, and failing to cope with these features of the data can lead to biased parameter estimates, and thus false conclusions and decisions.

Negative Binomial (NB) regression is widely considered as the default choice for data that are over-dispersed relative to Poisson. However, NB regression may not be the best choice for power-law data with long tails, or for highly skewed data with an excessive number of zeros due to the rare occurrence of non-zero events. The latter often requires the application of zero-inflated and hurdle models. In these models, two data generation processes are considered: one that generates zero counts and the other that generates the positive counts. These models can be complicated, with additional parameters to estimate. Thus, if one is not interested into distinguishing the two different processes for the data generation, it would be more appropriate to consider a simple model that can cope with skewed data with excessive zeros. In addition, NB regression cannot deal with data that are under-dispersed relative to Poisson. There have been some attempts to extend Poisson regression based models to under dispersion, such as the generalised Poisson regression model (Efron, 1986; Famoye, 1993), COM-Poisson regression (Sellers and Shmueli, 2010) or hyper-Poisson regression models (Sáez-Castillo and Conde-Sánchez, 2013). However, these models are all modifications of a Poisson model and have been shown to be rather complex and computationally expensive in practice.

This paper aims to introduce a simple count regression model, namely the Discrete Weibull (DW) regression model, and to show how this model can capture different levels of dispersion adaptively. This is a challenge faced by existing count regression models. Moreover, we show how the DW model can capture

power-law behaviour, excessive zeros or high skewness of the underlying distributions without the need for an additional mixture component.

The motivation behind considering the DW distribution (Nakagawa and Osaki, 1975) stems from the vital role played by the continuous Weibull distribution in the survival analysis and failure time studies. The estimation and inference for parameters of a DW distribution have been investigated in a small number of studies. Khan et al. (1989) proposed the method of proportion whereas Kulasekera (1994) suggested maximum likelihood estimators of the DW parameters based on type I censored samples. The count data application examples of DW include Englehardt and Li (2011) and Englehardt et al. (2012), who showed that the counts of living microbes (pathogen) in water are highly skewed and can be efficiently modelled using a DW distribution. Section 2 provides a review and description of the DW distribution and its properties. Section 3 illustrates the ability of a DW distribution to model data that are both over- and under-dispersed relative to Poisson. The DW regression model is introduced in Section 4 to investigate the relationship between a count response and a set of covariates. This model is applied to a number of real datasets in Section 6, showing the ability of the model to capture over- and under-dispersion, as well as excessive zeros and skewness.

2 Discrete Weibull distribution

If Y follows a (type 1) DW distribution (Nakagawa and Osaki, 1975), then the cumulative distribution function of Y is given by

$$F(y;q,\beta) = \begin{cases} 1-q^{(y+1)^\beta} & \text{for } y = 0,1,2,3,\ldots \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

and its probability mass function by

$$f(y;q,\beta) = \begin{cases} q^{y^\beta} - q^{(y+1)^\beta} & \text{for } y = 0,1,2,3,\ldots \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

with the parameters $0 < q < 1$ and $\beta > 0$.

The first two moments of a DW distribution are given by

$$\begin{aligned} E(Y) &= \mu = \sum_{y=1}^{\infty} q^{y^\beta} \\ E(Y^2) &= 2\sum_{y=1}^{\infty} y q^{y^\beta} - E(Y). \end{aligned} \tag{3}$$

From these, the variance for a DW distribution is given

$$\sigma^2 = 2\sum_{y=1}^{\infty} yq^{y^\beta} - \mu - \mu^2. \tag{4}$$

The moments can be calculated numerically using a truncated support (Barbiero, 2015).

A nice property of the DW distribution is that its τth ($0 < \tau < 1$) quantile, that is the smallest value of y for which $F(y) \geq \tau$, has a close form expression, given by

$$Q(\tau) = \left\lceil \left(\frac{\log(1-\tau)}{\log(q)} \right)^{\frac{1}{\beta}} - 1 \right\rceil. \tag{5}$$

for $\tau \geq 1 - q$. This is in contrast to Poisson and NB, which do not have a closed form expression for the quantiles.

2.1 Parameter estimation

Given a sample $y_1, y_2, \ldots, y_n$, from a DW distribution, the log-likelihood can be written as:

$$\ell = \sum_{i=1}^{n} \log\left(q^{y_i^\beta} - q^{(y_i+1)^\beta} \right),$$

from which the MLEs of q and β can be easily obtained by directly maximising this log-likelihood using any non-linear optimisation tool.

2.2 Properties and special cases of a DW distribution

- It can be seen from Equation 2 that,

  $$f(0) = 1 - q.$$

  Thus, when q is small, an excessive zero case occurs.

- The discrete Rayleigh distribution in Roy (2004) is a special case of a DW with $\beta = 2$ and $q = \theta$.

- The geometric distribution is a special case of a DW, with $\beta = 1$ and $q = 1 - p$. Moreover, for the geometric distribution, the variance is always greater than its mean. Therefore, a DW with $\beta = 1$ is a case of over-dispersion relative to Poisson, regardless of the value of q. In particular, when $\beta = 1$ and $q = e^{-\lambda}$, the distribution is the discrete exponential distribution introduced by Sato et al. (1999).

- β can be considered as controlling the range of values of the variable. As $\beta \to \infty$, the DW approaches a Bernoulli distribution with probability q.

In the next section, we discuss a property of DW that is particularly advantageous as a model for count data.

3 DW accounts for different types of dispersion

Dispersion in count data is formally defined in relation to a specified model being fitted to the data (Cameron and Trivedi, 2013; Hilbe, 2014). In particular, let

$$VR = \frac{\text{observed variance}}{\text{theoretical variance}}. \tag{6}$$

So VR is the ratio between the observed variance from the data and the theoretical variance from the model. Then the data are said to be over-/equi-/under- dispersed relative to the fitted model if the observed variance is larger/equal/smaller than the theoretical variance specified by the model, respectively. It is common to refer to dispersion relative to Poisson. In that case, the variance of the model is estimated by the sample mean. Thus, over-/equi-/under- dispersion relative to Poisson refers to cases where the sample variance is larger/equal/smaller than the sample mean, respectively. Since the theoretical variance of a NB is always greater than its mean, as $\sigma^2 = \mu + \frac{1}{k}\mu^2$ for $k > 0$, the NB regression model is the natural choice for data that are over-dispersed relative to Poisson. However, crucially, NB cannot handle under-dispersed data.

In contrast to this, we show how a DW distribution can handle data that are both over- and under- dispersed relative to Poisson. In particular, Figure 1 shows the VR values in Equation 6 for data simulated by a $DW(\beta, 0.7)$ and fitted by Poisson and NB distributions. The plots show cases of data over and under dispersed relative to Poisson. In addition, while NB can fit well data that are over-dispersed relative to Poisson (i.e. VR close to 1), this does not happen for under-dispersed data where both Poisson and NB are inappropriate. Figure 2 considers more closely the case of dispersion relative to Poisson for a range of values of q and β and shows how DW, a single distribution with as many parameters as NB, can capture both cases of under-, equi- and over- dispersion relative to Poisson. In particular, the plot shows that:

- $0 < \beta \leq 1$ is a case of over-dispersion, regardless of the value of q.
- $\beta \geq 2$ is a case of under-dispersion, regardless of the value of q. In fact, DW approaches the Bernoulli distribution with mean p and variance $p(1-p)$ for $\beta \to \infty$.

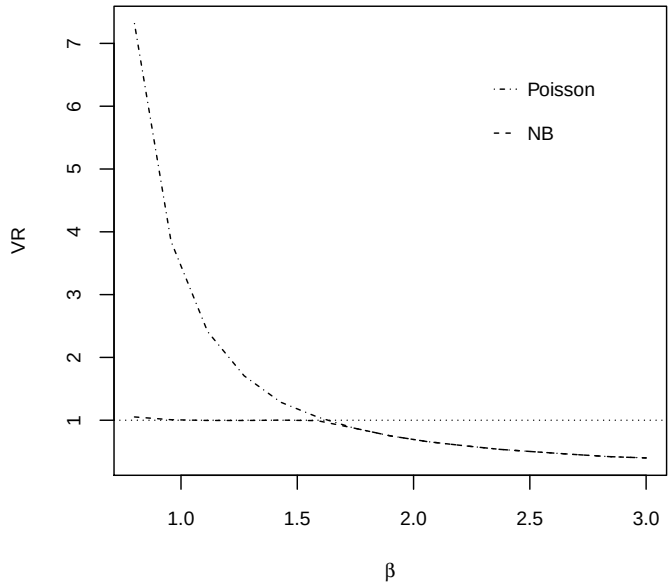

Figure 1: Ratio of observed and theoretical variance of data simulated by $DW(0.7, \beta)$ and fitted by Poisson and NB.

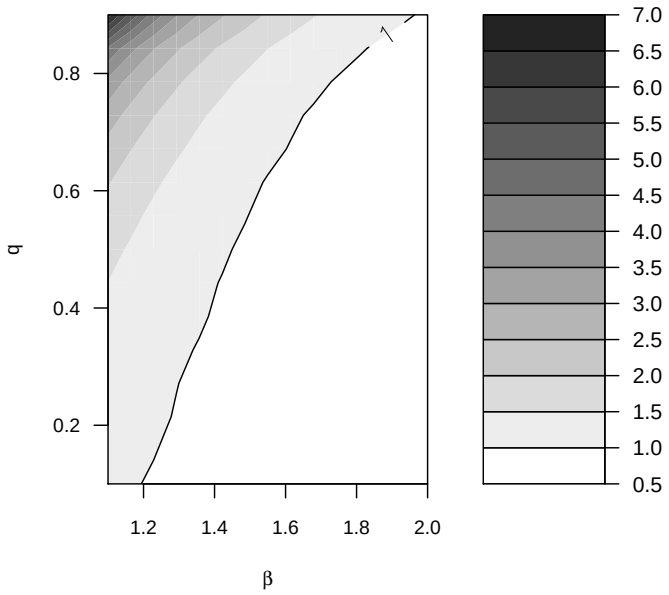

Figure 2: Ratio of observed and theoretical variance of data simulated by $DW(q,\beta)$ and fitted by Poisson. The white area corresponds to data under-dispersed relative to Poisson whereas the grey shaded area corresponds to over-dispersion.

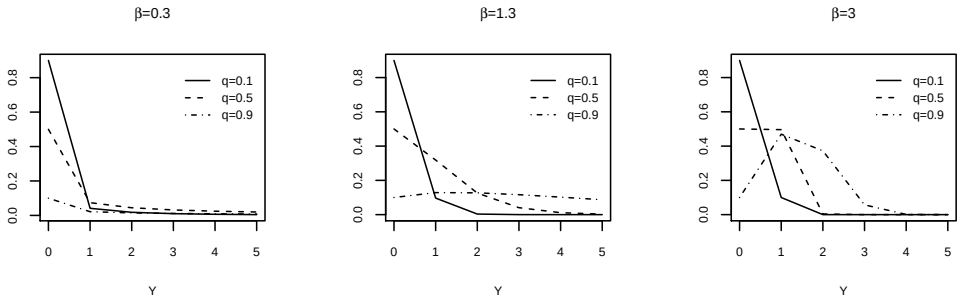

Figure 3: Effect of q on the probability mass function of the DW for different βs.

- $1 < \beta < 2$ leads to both cases of over and under-dispersion depending on the value of q.

4 DW regression model

We now exploit this advantageous property of a DW distribution within a regression context, where the interest is to model the relationship between a count response variable and a set of covariates.

4.1 Model formulation

We introduce the DW regression model for count data in analogy with the continuous Weibull regression, which is well-known in survival analysis and life time modelling. Recalling that the distribution function of a continuous Weibull distribution is given by

$$F(y;\lambda,\beta) = 1 - e^{-\lambda y^{\beta}}, y \geq 0,$$

with scale parameter λ, one can see that the parameter q of a DW distribution is equivalent to $e^{-\lambda}$ in the continuous case. Since Weibull regression imposes a log link between the parameter λ and the predictors (Da Silva et al., 2008; Lee and Wang, 2003), the DW regression can be introduced via the parameter q. Figure 3 shows how the parameter q affects the scale and the shape of the probability mass function of the DW distribution.

Therefore, in order to introduce a DW regression model, we assume that, for $i = 1, 2, \ldots, n$, the response Y_i has a DW conditional distribution $f(y_i, q(x_i), \beta | x_i)$, where $q(x_i)$ is the DW parameter related to the explanatory variables x_i through the link function

$$\log(-\log(q_i)) = \boldsymbol{x}_i'\boldsymbol{\alpha}, \qquad \boldsymbol{x}_i'\boldsymbol{\alpha} = \alpha_0 + x_{i1}\alpha_1 + \ldots + x_{iP}\alpha_P. \tag{7}$$

This transforms q from the probability scale (i.e the interval $[0,1]$) to the interval $[-\infty,+\infty]$ and ensures that the parameter q remains in $[0,1]$. Indeed, from Equation 7, q_i can be expressed as:

$$q_i = e^{-e^{\boldsymbol{x}_i'\boldsymbol{\alpha}}}, \tag{8}$$

from which the conditional probability mass function of the response variable Y_i given x_i is as follows:

$$f(y_i|x_i) = \left(e^{-e^{\boldsymbol{x}_i'\boldsymbol{\alpha}}}\right)^{y_i^{\beta}} - \left(e^{-e^{\boldsymbol{x}_i'\boldsymbol{\alpha}}}\right)^{(y_i+1)^{\beta}}. \tag{9}$$

Finally, in order to obtain the MLEs for the unknown parameters $\boldsymbol{\alpha}$ and β, the log-likelihood of Equation 9 is maximised numerically using standard optimisation tools.

4.2 Interpretation of the Regression Coefficients

After a DW regression model has been estimated, the following can be obtained:

- The fitted values for the central trend of the conditional distribution can be obtained in one of two methods:
  - mean: the Equation 3, as mentioned earlier, can be calculated numerically using the approximated moments of the DW (Barbiero, 2015).
  - median: the quantile formula provided in Equation 5, can be applied. Due to the skewness, which is common for count data, the median is more appropriate than the mean. The fitted conditional median can be obtained easily from the closed form expression of quantiles for DW, as
$$M(\boldsymbol{x}) = \left\lceil \left(-\frac{\log(2)}{\log(q(\boldsymbol{x}))}\right)^{\frac{1}{\beta}} - 1 \right\rceil. \tag{10}$$
- The conditional quantile for any τ can be obtained from Equation 5.

By substituting Equation 8 in Equation 10, we have that

$$\log\left(M(\boldsymbol{x})+1\right) = \frac{1}{\beta}\log\big(\log(2)\big) - \frac{1}{\beta}\boldsymbol{x}'\boldsymbol{\alpha}. \tag{11}$$

Thus, the regression parameters $\boldsymbol{\alpha}$ can be interpreted in relation to the log of the median. This is in analogy with Poisson and NB models, where the parameters are

linked to the mean. In particular, $\frac{\log\big(\log(2)\big)-\alpha_0}{\beta}$ is related to the conditional median when all covariates are set to zero, whereas $\frac{-\alpha_p}{\beta}$, $p=1,\ldots,P$, can be related to the change in the median of the response corresponding to a one unit change of X_p, keeping all other covariates constant.

4.3 Diagnostics Checks

After fitting a DW regression model, it is essential to consider a diagnostics' analysis to investigate the appropriateness of the model. Given that the response is discrete, we advise performing a residual analysis based on the randomized quantile residuals, as developed by (Dunn and Smyth, 1996) and used in many other studies, e.g. (Ospina and Ferrari, 2012; Vanegas et al., 2013). In particular, let

$$r_i = \Phi^{-1}(u_i) \qquad , \qquad i=1,\ldots,n \tag{12}$$

where $\Phi(.)$ is the standard normal distribution function and u_i is a uniform random variable on the interval

$$\begin{aligned}(a_i,b_i] &= \left(\lim_{y\uparrow y_i} F(y;\hat{q}_i,\hat{\beta}), F(y_i;\hat{q}_i,\hat{\beta})\right] \\ &\approx \left[F(y_i-1;\hat{q}_i,\hat{\beta}), F(y_i;\hat{q}_i,\hat{\beta})\right].\end{aligned}$$

These residuals follow the standard normal distribution apart from sampling variability in $\hat{q}_i$ and $\hat{\beta}$. Hence, the validity of a DW model can be assessed using goodness-of-fit investigations of the normality of the residuals, such as Q-Q plots and normality tests. Simulated envelopes can be added to the Q-Q plots, as in (Ferrari and Cribari-Neto, 2004; Garay et al., 2011; Sáez-Castillo and Conde-Sánchez, 2013; Atkinson, 1985).

In addition to the residual analysis, it is informative also to check whether the data shows any under or over dispersion relative to the specified DW conditional distribution. In the case of good fitting, we would expect the ratio of observed and theoretical variance in equation 6 to be close to 1 for each x. We will expand more on this point in the next section.

5 DW regression naturally handles covariate-specific dispersion

We have already shown in Section 3 how DW can model data that are underdispersed relative to Poisson. In this section, we would like to investigate this

Table 1: Simulation study: DW parameter estimates by Maximum Likelihood (MLE), together with bias, Mean Squared Error (MSE) and length of the 95% Confidence Interval, averaged over 1000 iterations.

	MLE	Bias	MSE	95% CI Length
α_0	0.5103	0.0103	0.0186	0.5191
α_1	0.4105	0.01046	0.0046	0.2556
α_2	-0.3038	-0.0038	0.0008	0.1074
β	1.6189	0.0189	0.0092	0.3586

further within a regression context. Here, it is also possible that the conditional variance is larger than the conditional mean for a specific covariate pattern (over-dispersion), but the conditional variance is smaller than the conditional mean for another covariate pattern (under-dispersion). In the literature, regression models for count data that can capture under-dispersion or both types of over- and under-dispersion simultaneously are in the form of extended versions of Poisson regression, such as quasi-Poisson, COM-Poisson or hyper-poisson (Sáez-Castillo and Conde-Sánchez, 2013). In the case of mixed types of dispersion, the dispersion parameter can be assumed to be linked to the covariates. However, a covariate-dependent dispersion increases the complexity of the model significantly and reduces its interpretability. So, in practice, most implementations fix the dispersion parameter and assume that only the mean is linked to the covariates. As the DW distribution naturally accounts for over and under-dispersion, a DW regression model becomes a simple and attractive alternative to existing regression models for count data.

We emphasize this point by a simple simulation study. We consider a case with two covariates, $X_1 \sim N(0,1)$ and $X_2 \sim \text{Uniform}(0,10)$. We assume the regression parameters to take values $\boldsymbol{\alpha} = (\alpha_0, \alpha_1, \alpha_2) = (0.5, 0.4, -0.3)$. In addition, the parameter β of the DW is assumed to be $\beta = 1.6$. Then, we sample 300 values of the covariates and the corresponding response from $DW(q_i, \beta)$, where q_i is calculated as in Equation 8, for $i = 1, \ldots, n$. Table 1 reports the estimates of the parameters, together with the average bias and the mean squared error (MSE) over 1000 iterations.

Figure 4 shows a plot of the dispersions in equation 6 in the case of Poisson, NB and DW fitting. Some note is required on the calculation of these ratios, as the observed variance cannot be computed for each individual covariate vector x. For the calculation, we split the response values into 10 groups of similar size, based on the percentiles of the linear predictors $\boldsymbol{x}'\boldsymbol{\alpha}$. Then the observed variance is computed within each group, while the theoretical variance is averaged within each group. If the model is well specified, we would expect these values to be

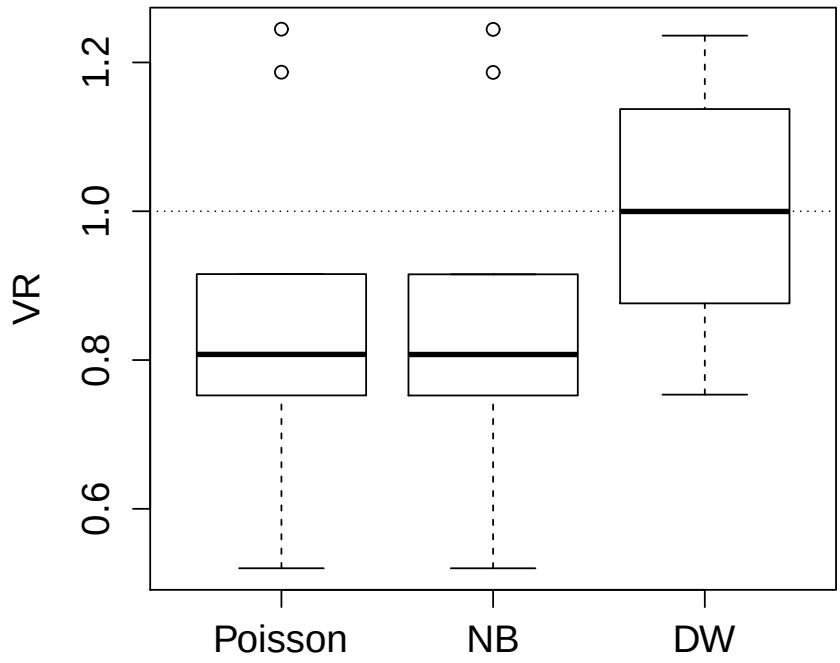

Figure 4: Distribution of ratios of observed and theoretical conditional variance on simulated data from the DW regression model, with the theoretical variance fitted by Poisson, NB and DW, respectively.

close to 1. This is shown in Figure 4 for DW which is the model used in the simulation. Poisson and NB show both under-dispersion in most cases and over-dispersion in two cases. Thus, this simulation shows a simple scenario of a mixed level of dispersion, which cannot be captured well by standard Poisson and NB models.

6 Application to real data sets

To demonstrate the ability of the DW regression model to handle over- and under-dispersion automatically, in this section DW regression is applied to different data sets that show various types of dispersions relative to Poisson. The first sub-

section includes an under-dispersed dataset, while the second includes an over-dispersed case. The third subsection focuses on a zero-inflated dataset. Finally, an illustrative example for the mixed level of dispersion is provided. Various popular count data regression models, namely Poisson regression (R function *glm*), NB regression (R function *glm.nb*), COM-Poisson regression (R package *COMPoissonReg* (Sellers and Lotze, 2011)), zero-inflated and hurdle models (*pscl* R package (Zeileis et al., 2008)), are applied and compared with DW regression by means of classical AIC and BIC criteria (Dayton, 2003).

6.1 The case of under-dispersion: inhaler usage data

For this example, we use data from Grunwald et al. (2011). They consist of 5209 observations and report the daily count of using (albuterol) asthma inhalers for 48 children suffering from asthma, aged between 6 and 13 during the school day, for a period of time at the Kunsberg School at National Jewish Health in Denver, Colorado. The main objective of this analysis is to investigate the relationship between the inhaler use (representing the asthma severity), and air pollution, which is recorded by four covariates: the percentage of humidity, the barometric pressure (in mmHG/1000), the average daily temperature (in Fahrenheit degree/100) and the morning levels of PM_{25}, which are small air particles less than 25mm in diameter. The response variable, which is the inhaler use count, has a sample mean of 1.2705 and variance of 0.8433, thus pointing to a case of under-dispersion relative to Poisson.

The results in Table 2 suggest that DW and COM-Poisson regression provide better fitting than both Poisson and NB models, according to both AIC and BIC. For DW, the parameters are reported with the parametrization linked to the median as previously described. The left panel of Figure 5 indicates under-dispersion relative to Poisson and NB across the full range of the covariates, and a good fit of DW compared to the other models (VR values close to 1). COM-Poisson could not be added to this plot due to the complexity in calculating the theoretical variances in this case. The right panel compares the observed and expected frequencies for the three models and shows again a good fit for DW. Finally, Figure 6 plots the randomized quantile residuals from the DW regression model, which are only moderately departing from normality (p-value of Kolmogorov-Smirnov (KS) test: 0.025).

6.2 The case of over-dispersion: strikes data

This dataset is available in the *Ecdat* R package (Croissant, 2015), under the name of *StrikeNb*. The response variable is the number of contract strikes in U.S. manufacturing observed monthly from January 1968 to December 1976. The predictor

Table 2: Maximum likelihood estimates, AIC and BIC from different regression models fitted to the inhaler use data.

	humidity	pressure	temperature	particles	other	AIC	BIC
Poisson	-0.1125	4.0950	-0.2035	0.0225	-	13915.47	13948.26
NB	-0.1125	4.0950	-0.2035	0.0225	$\hat{k}$=31905.28	13917.54	13956.89
COM-Poisson	-0.1727	6.2838	-0.3127	0.0347	$\hat{\nu}$=1.9205	13450.77	13490.12
DW	-0.1050	2.6376	-0.1735	0.0136	$\hat{\beta}$=2.1277	13484.36	13523.71

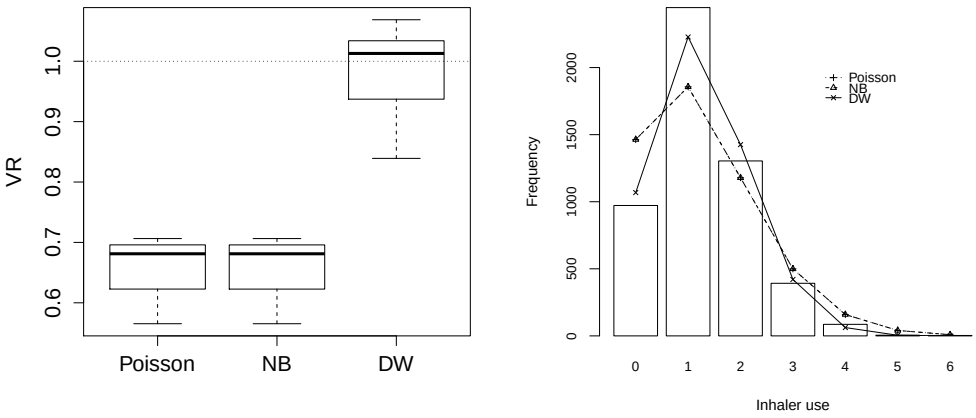

Figure 5: Comparison of DW with NB and Poisson on inhaler use data. Left: Distribution of ratios of observed and theoretical conditional variance on the data fitted by Poisson, NB and DW, respectively. Right: Observed and expected frequencies for each model.

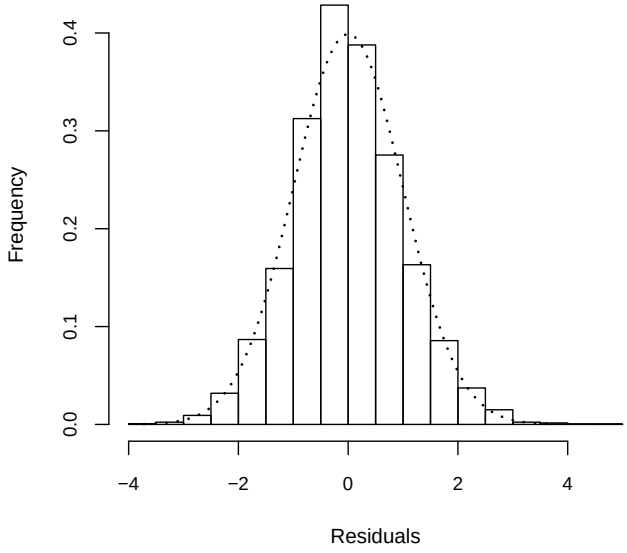

Figure 6: Residuals analysis for inhaler use data for the DW regression model: histogram of randomized quantile residuals with superimposed N(0,1) density (dotted line).

Table 3: Maximum likelihood estimates, AIC and BIC from different regression models fitted to the strikes data.

	economic activity	other	AIC	BIC
Poisson	3.1342	-	627.9689	633.3332
NB	3.2250	$\hat{k}$=3.1849	566.5969	574.6433
DW	3.2043	$\hat{\beta}$=1.6527	564.157	572.2034

is the level of economic activity, which is measured as cyclical departure of aggregate production from its trend level. The response variable has a sample mean of 5.2407 and variance of 14.0723, suggesting over-dispersion relative to Poisson. Indeed, a comparison of Poisson and NB distributions solely on the response variable using a likelihood ratio test (*lmtest* R package, (Zeileis and Hothorn, 2002)) shows evidence of over-dispersion with a chi-square test statistic of 73.024 and a p-value of < 0.001.

After fitting three regression models and comparing them via AIC and BIC, Table 3 shows that the DW model is only marginally superior to NB, but both DW and NB give much better fit to the data than the Poisson regression model. The left panel of Figure 7, indicates a case of over-dispersion relative to Poisson across the whole range of the covariates. Additionally, it indicates the good fitting by NB and DW with VR values closer to 1. The right panel confirms the good fit of NB and DW. For visualization purposes, the small number of observations larger than 11 are grouped together on this plot. Finally, Figure 8 shows that the residuals follow closely a normal distribution (KS p-value 0.83), with not many points falling outside the simulated 95% envelope's bounds.

6.3 The case of excessive zeros: doctor visits data

The following dataset illustrates the case of excessive zero counts. Thus, besides the Poisson, NB, and DW regression, we will include also zero-inflated and hurdle models in the comparison. For these, we consider the logit link function for the binomial distribution representing the probability of the extra zeros (R package *pscl*, (Zeileis et al., 2008)). The data are available from the *Ecdat* R package, under the name *Doctor*. The data consist of 485 observations from the United States in the year 1986, and contain four variables for each patient: the number of doctor visits, which is taken as response, the number of children in the household, a measure of access to healthcare and a measure of health status (larger positive numbers are associated with poorer health). The response variable in this study, the number of doctor visits, has approximately 50% of zeros, and thus it can be considered as a case of excessive zeros. Indeed, the response variable has a mean

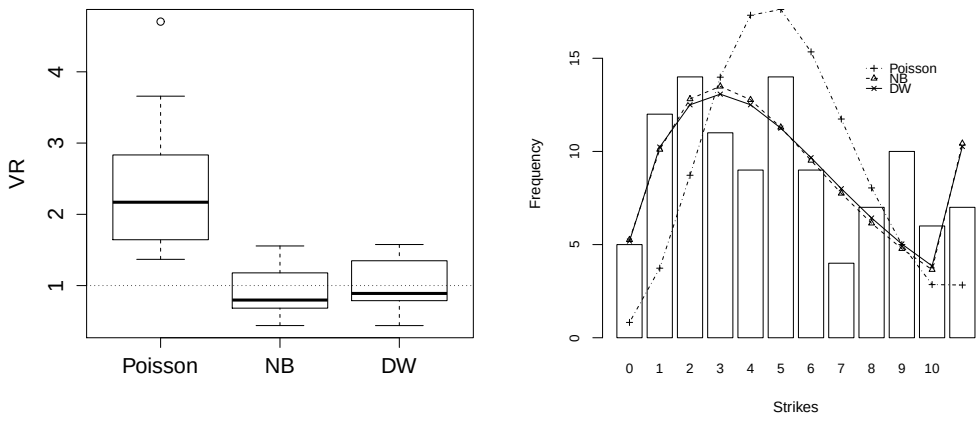

Figure 7: Comparison of DW with NB and Poisson on strikes data. Left: Distribution of ratios of observed and theoretical conditional variance on the data fitted by Poisson, NB and DW, respectively. Right: Observed and expected frequencies for each model.

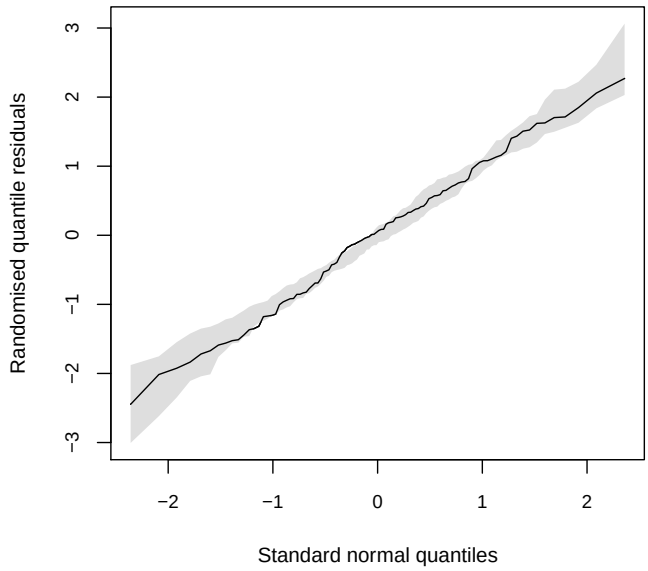

Figure 8: Residuals analysis for strikes data: Q-Q plot of randomized quantile residuals of the DW regression model.

Table 4: Maximum likelihood estimates, AIC and BIC from different regression models fitted to doctor visits data

	children	access	health	other	AIC	BIC
Poisson	-0.1759	0.9369	0.2898	-	2179.487	2196.223
NB	-0.1706	0.4197	0.3154	$\hat{k}$=0.5525	1581.88	1602.801
Zero-inflated Models						
Poisson						
count model	-0.1498	0.8053	0.1736	-	1885.813	1919.287
logit model	0.0843	-0.1048	-0.4147	-		
NB						
count model	-0.1414	0.6491	0.2239	$\hat{k}$=0.6869	1578.5	1616.158
logit model	0.2465	1.2085	-2.0676	-		
Hurdle Models						
logit model	-0.1462	0.4252	0.4524	-	-	-
Poisson count model	-0.1506	0.8143	0.1733	-	1885.808	1919.281
NB count model	-0.1664	0.5404	0.2157	$\hat{k}$=0.2596	1576.302	1613.959
DW	-0.1309	0.34029	0.2758	$\hat{\beta}$=0.7823	1575.796	1596.717

of 1.6103 and variance of 11.2011 and the likelihood ratio test between NB and Poisson return a test statistic of 804.59 and a p-value of < 0.001.

Table 4 shows the best fit for the DW regression model in terms of AIC and BIC. Figure 9 shows a case of over-dispersion relative to Poisson across the full range of covariates and a good fit for DW and ZINB. We exclude Poisson from the plot for visualization purposes, as the VR values are very large in this case. As in the previous example, the residuals of the DW model are well approximated by a normal distribution (KS p-value 0.9239). This example shows how DW can model also cases of excessive zeros, without the need of additional parameters as in the zero-inflated models.

6.4 The case of a mixed level of dispersion: bids data

In this section, we report the analysis of a dataset where a mixed level of dispersion was observed, that is the conditional distribution is over-dispersed relative to Poisson for some covariate pattern but is under-dispersed for another covariate pattern. The data are taken from Cameron and Johansson (1997) and are available in the *Ecdat* R package under the name of *Bids*. The data record the number of bids received by 126 US firms that were targets of tender offers during a certain period of time. The dependent variable here is the number of bids, with a mean of 1.7381 and a variance of 2.0509. The objective of the study is to investigate the effect of some variables on the number of bids. For this analysis, we consider the following covariates: bid price, taken as the price at a particular week divided by the price 14 working days before the bid, size, that is the total book value of

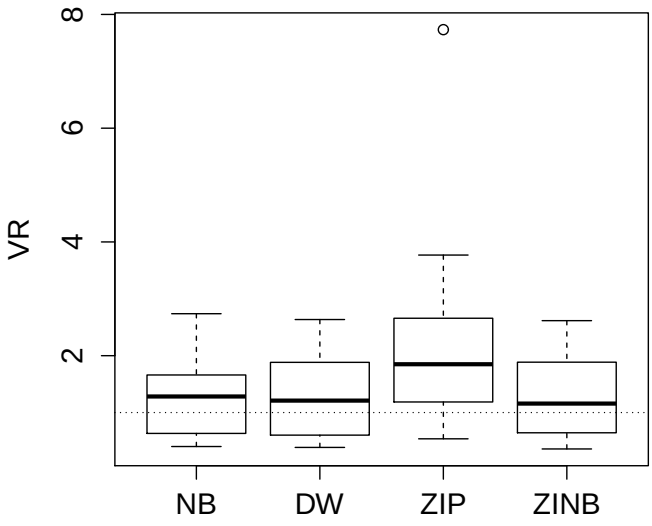

Figure 9: Distribution of ratios of observed and theoretical conditional variance on the doctor visits data fitted by Poisson, NB and DW, respectively.

Table 5: Maximum likelihood estimates, AIC and BIC from different regression models fitted to bids data.

	price	size	regulator	other	AIC	BIC
Poisson	-0.7849	0.0362	0.0547	-	402.2602	413.6054
NB	-0.7824	0.0369	0.0544	$\hat{k}$=33.3289	403.9481	418.1295
DW	-0.6761	0.0552	0.0293	$\hat{\beta}$=1.9403	395.1214	409.3028

assets measured in billions dollars, and regulator, a dummy variable which is 1 if there was an intervention by federal regulators and 0 otherwise.

Figure 10 and Table 5 shows once again a very good fit of the DW regression model to these data, compared to Poisson and NB. Figure 10 shows a mixed level of dispersion relative to Poisson and NB, with most covariate patterns leading to under-dispersion, but with a small number of over-dispersed cases. The DW model has a clearer distribution of VR values around 1 and it also fits the data well, with a KS p-value of 0.11 for the randomized quantile residuals.

7 Conclusion

In this paper we introduce a regression model based on a DW distribution and show how this model can be seen as a simple and unified model to capture different levels of dispersion in the data, namely under-dispersion and over-dispersion relative to Poisson, including the common case of excessive zeros. This is an attractive feature of DW, similarly to the flexibility of the continuous Weibull distribution in adapting to a variety of hazard rates. In addition, the proposed DW regression model, unlike generalised linear models in which the conditional mean is central to the interpretation, has the advantage of modeling the whole conditional density, including all conditional quantiles which can be easily extracted from the fitted model. This is particularly useful as most count data are highly skewed, even power-law.

A popular model for under-dispersion is the COM-Poisson regression model. However, the probability mass function of COM-Poisson is not in a closed form and contains an infinite sum, which requires an approximate computation. In fact, the COM-Poisson implementation which was used for the examples in this paper required more computational time than the DW regression model, which uses a straightforward maximum likelihood estimation procedure. This is particularly beneficial in the case of large sample sizes. While NB is the most applied model for over-dispersion, the DW regression model is shown to be an attractive alternative to the NB regression model for over-dispersion. In particular, several examples in this study show that DW regression provides the best fitting model,

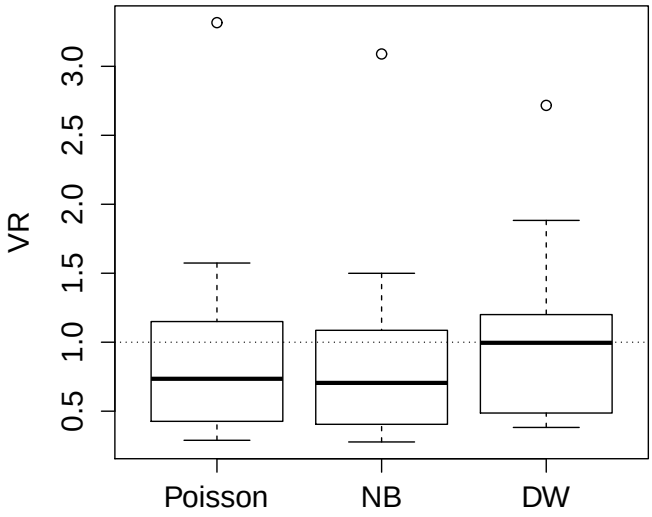

Figure 10: Distribution of ratios of observed and theoretical conditional variance on the bids data fitted by Poisson, NB and DW, respectively.

both in cases of over- and under-dispersion, and is also able to capture situations of a mixed-level of dispersion. In addition, the DW regression model can be applied to data with an excessive number of zero counts without requiring the need for additional parameters, as in the case of zero-inflated or hurdle models.

The DW regression model described in this paper is implemented in the R package *DWreg*, freely available in CRAN,
(`http://CRAN.R-project.org/package=DWreg`).

References

Atkinson, A. C., 1985. Plots, transformations, and regression: an introduction to graphical methods of diagnostic regression analysis. Clarendon Press Oxford.

Barbiero, A., 2015. DiscreteWeibull: Discrete Weibull Distributions (Type 1 and 3). R package version 1.0.1.
URL `http://CRAN.R-project.org/package=DiscreteWeibull`

Cameron, A. C., Johansson, P., 1997. Count data regression using series expansions: with applications. Journal of Applied Econometrics 12, 203–223.

Cameron, A. C., Trivedi, P. K., 2013. Regression analysis of count data. Cambridge University Press.

Croissant, Y., 2015. Ecdat: Data Sets for Econometrics. R package version 0.2-9.
URL `http://CRAN.R-project.org/package=Ecdat`

Da Silva, M. F., Ferrari, S. L., Cribari-Neto, F., 2008. Improved likelihood inference for the shape parameter in Weibull regression. Journal of Statistical Computation and Simulation 78 (9), 789–811.

Dayton, C. M., 2003. Model comparisons using information measures. Journal of Modern Applied Statistical Methods 2 (2), 2.

Dunn, P. K., Smyth, G. K., 1996. Randomized quantile residuals. Journal of Computational and Graphical Statistics 5 (3), 236–244.

Efron, B., 1986. Double exponential families and their use in generalized linear regression. Journal of the American Statistical Association 81 (395), 709–721.

Englehardt, J. D., Ashbolt, N. J., Loewenstine, C., Gadzinski, E. R., Ayenu-Prah, A. Y., 2012. Methods for assessing long-term mean pathogen count in drinking water and risk management implications. Journal of Water and Health 10 (2), 197–208.

Englehardt, J. D., Li, R., 2011. The discrete Weibull distribution: an alternative for correlated counts with confirmation for microbial counts in water. Risk Analysis 31 (3), 370–381.

Famoye, F., 1993. Restricted generalized Poisson regression model. Communications in Statistics-Theory and Methods 22 (5), 1335–1354.

Ferrari, S., Cribari-Neto, F., 2004. Beta regression for modelling rates and proportions. Journal of Applied Statistics 31 (7), 799–815.

Garay, A. M., Hashimoto, E. M., Ortega, E. M., Lachos, V. H., 2011. On estimation and influence diagnostics for zero-inflated negative binomial regression models. Computational Statistics & Data Analysis 55 (3), 1304–1318.

Grunwald, G. K., Bruce, S. L., Jiang, L., Strand, M., Rabinovitch, N., 2011. A statistical model for under-or overdispersed clustered and longitudinal count data. Biometrical Journal 53 (4), 578–594.

Hilbe, J. M., 2014. Modeling Count Data. Cambridge University Press.

Khan, M. A., Khalique, A., Abouammoh, A., 1989. On estimating parameters in a discrete Weibull distribution. IEEE Transactions on Reliability 38 (3), 348–350.

Kulasekera, K., 1994. Approximate MLE's of the parameters of a discrete Weibull distribution with type i censored data. Microelectronics Reliability 34 (7), 1185–1188.

Lee, E. T., Wang, J., 2003. Statistical methods for survival data analysis. Vol. 476. John Wiley & Sons.

Nakagawa, T., Osaki, S., 1975. The discrete Weibull distribution. IEEE Transactions on Reliability 24 (5), 300–301.

Nelder, J. A., Wedderburn, R. W., 1972. Generalized linear models. Journal of the Royal Statistical Society. Series A 135, 370–384.

Ospina, R., Ferrari, S. L., 2012. A general class of zero-or-one inflated beta regression models. Computational Statistics & Data Analysis 56 (6), 1609–1623.

Roy, D., 2004. Discrete Rayleigh distribution. IEEE Transactions on Reliability 53 (2), 255–260.

Sáez-Castillo, A., Conde-Sánchez, A., 2013. A hyper-Poisson regression model for overdispersed and underdispersed count data. Computational Statistics & Data Analysis 61, 148–157.

Sato, H., Ikota, M., Sugimoto, A., Masuda, H., 1999. A new defect distribution metrology with a consistent discrete exponential formula and its applications. IEEE Transactions on Semiconductor Manufacturing 12 (4), 409–418.

Sellers, K., Lotze, T., 2011. COMPoissonReg: Conway-Maxwell Poisson (COM-Poisson) Regression. R package version 0.3.4.
URL http://CRAN.R-project.org/package=COMPoissonReg

Sellers, K. F., Shmueli, G., 2010. A flexible regression model for count data. Annals of Applied Statistics 4 (2), 943–961.

Vanegas, L. H., Rondón, L. M., Cordeiro, G. M., 2013. Diagnostic tools in generalized Weibull linear regression models. Journal of Statistical Computation and Simulation 83 (12), 2315–2338.

Zeileis, A., Hothorn, T., 2002. Diagnostic checking in regression relationships. R News 2 (3), 7–10.
URL http://CRAN.R-project.org/doc/Rnews/

Zeileis, A., Kleiber, C., Jackman, S., 2008. Regression models for count data in r. Journal of Statistical Software 27.